

**APPLICATION
FOR
UNITED STATES LETTERS PATENT**

APPLICANT NAME: MARY ELLEN SIKSA

**TITLE: AUTOMATIC BILLING EVENT SUBMISSION
RECONCILIATION FOR ON DEMAND SYSTEMS**

DOCKET NO.:END920030154US1

INTERNATIONAL BUSINESS MACHINES CORPORATION

AUTOMATIC BILLING EVENT SUBMISSION RECONCILIATION FOR ON DEMAND SYSTEMS

Mary Ellen Siksa

5

FIELD OF THE INVENTION

The invention described herein relates to a method, system, and program product for recording, analyzing, verifying, and reporting of product used, sold, or transferred from multiple locations at various prices or costs to different customers in a business or in commerce, and generating consolidated billing notices, for example notices of payment due, or funds transferred, in response to the transaction. Product can include, strictly by way of example telecommunications bandwidth, electricity, fuels, chemicals, and general merchandise.

15

BACKGROUND

Today service providers are beginning to embrace a pay-as-you-go billing model. This is a model where the user is billed for the amount of service used or resource consumed rather than a fixed, flat monthly fee. However, in the pay-as-you-go billing model, the service provider must be able to meter, monitor, or measure the customer's usage of a resource or service. Additionally, where necessary, the resource or service provider must be able to identify the customer's usage of the resource or commodity to both the class (or quality) of the resource or commodity, and to the quantity of the resource or commodity.

25

Consider, for example, cellular telephone service. In a typical cellular telephone service plan the customer is billed by the number of on air minutes. The number of on air

END920030154US1

minutes is then sent to a billing system where it is applied against the customer's plan, for example, \$29.95 per month for the first 300 on air minutes and thereafter \$0.25 per minute for each minute over 300 minutes.

- 5 Consider a more complex system, residential electrical service may be delivered through "multiple meters." One meter is for "base line" service at, for example \$0.11 per kilowatt hour for the first 600 kilowatt hours, and an increasing amount for each additional 200 kilowatt hours. Another meter is for "interruptible" power at \$0.09 per kilowatt hours. A third meter may be for "time shifted" power at a high rate during daytime hours and a
- 10 reduced rate for evening and night time use. Finally, there may be a "backwards" meter for "selling back" co-generated power. Consumption of the different classes of service is sent to a billing system where it is applied against the customer's plan and billed to the customer.
- 15 In a pay-as-you-go billing model, numerous events are measured in a distributed infrastructure, such as a cellular telephone or electric power service area, and are sent to a central billing system to charge the customer for the services or resources. One requirement of a billing system in a distributed pay-as-you-go billing system is accountability and tracking. There must be a way to ensure that billable events generated
- 20 by feeders at the point of use or consumption are received at the billing system. This requires reconciliation of feeder's feeds to the billing system with the feeds actually received by the billing system.

Figure 1, captioned as "PRIOR ART" illustrates the flow of billing events from a feeder

25 to a billing system. The feeder logs the records sent to the billing system in a feeder log, and the billing system logs the records received by the billing system in a billing log. To be noted is that the reports generated by the feeder and logged in the feeder log may not match the events received at the billing system and logged in the billing log.

END920030154US1

Currently the reconciliation between feeder records and billing records is performed manually, that is, by comparing reports generated at the many feeders with reports actually received and entered at the billing system. As pay-as-you-go billing systems become prevalent for service and resource providers, the sheer volumes of records flowing from feeders to the billing systems will increase, and manual systems, even “management by exception” manual systems, will no longer be acceptable. This means that there is a clear need for a reliable, automated way to reconcile feeder records to billing system records.

10

SUMMARY OF THE INVENTION

The method, system, and program product described herein provides a reliable, automated way to reconcile feeder records to billing system records. Reconciliation may be accomplished in real time or in a batch process. The method, system, and program product further provide notification when the feed records and billing records are out of balance. A further aspect of the method, system, and program product is the capability to scan the various log files on a random or scheduled or as-needed basis, examining the log files for missing, duplicate, and corrupted records, and generate balance reports and out of balance notifications, as necessary.

20

Aspects of the method, system, and program product include a mechanism to uniquely identify feeders, target billing systems, and even individual records. A further aspect of the method, system, and program product is the capability of handling both batch and real time feeds independent of the data format or communications protocol. A still further aspect of the method, system, and program product described herein is the handling of multiple feeders, multiple independent hops, and multiple target billing systems. A still

25

END920030154US1

further aspect of the method, system, and program product described herein is reporting capability to show in-balance and out-of-balance conditions.

According to the method, system, and program product described herein, an administrator
5 is notified of any error condition, including information for problem determination. The information may include a list of records sent by the various feeders and not received by the billing system, as well as information received by the billing system but neither identified to or associated to a feeder. This allows determination as to where the records were lost or duplicated.

10

THE FIGURES

Various aspects of the method, system, and program product described herein are illustrated in the FIGURES attached hereto.

15

FIGURE 1, denominated "Prior Art" illustrates a system of the prior art with a single feeder and single billing system, each with a log.

FIGURE 2 is a simplified flow chart showing data collection, data reconciliation, and
20 charge data consolidation (that is, billing).

FIGURE 3 is a flow chart of the reconciliation process. The reconciliation starts by logging the session start time and retrieving the last session start time, and retrieving the logs from the feeder and the interim control points. If there is an error at this point, the
25 appropriate notification is issued, the session status is logged, and the reconciliation process ended. If there is no error, the records are compared, checked for unreconciled records, and the report prepared and published. Session status is logged, and the reconciliation process is ended.

END920030154US1

FIGURE 4 illustrates a simplified system with a single feeder to a single billing system.

The feeder is associated to a feeder log, and to an interim node with an associated interim log. The feeder side of the system is separated from the billing side by a fire wall. The

5 billing side includes an interim node with an interim node log, and the billing system with an associated billing log.

FIGURE 5 illustrates a complex system with multiple feeders to a single billing system.

The feeders are associated to feeder logs, and to an interim node with an associated

10 interim log. The feeder side of the system is separated from the billing side by a fire wall.

The billing side includes an interim node with an interim node log, and the billing system with an associated billing log.

FIGURE 6 illustrates a still more complex system with multiple feeders to multiple

15 billing systems. The feeders are associated to feeder logs, and to an interim node with an associated interim log. The feeder side of the system is separated from the billing side by a fire wall. The billing side includes an interim node with an interim node log, and the billing systems with associated billing logs.

20 FIGURE 7 illustrates a data collection – data reconciliation – billing system applied to an electric power distribution system. The data collection – data reconciliation – billing

system has multiple feeders, typically “electric meters” to multiple billing systems. The

feeders collect data from the various associated “electric meters” and are associated to

feeder logs, and to an interim node with an associated interim log. The feeder side of the

25 system is separated from the billing side by a fire wall. The billing side includes an

interim node with an interim node log, and the billing systems with associated billing

logs. In a purely deregulated environment the different billing systems could be different generating companies, such as a coal fired company, a gas fired company, a hydroelectric

END920030154US1

company, and a wind turbine company. Alternatively, the different billing systems could represent different markets or regulatory regimes.

DETAILED DESCRIPTION

5

The method, system, and program product described herein provides a reliable, automated way to reconcile feeder records to billing system records. Reconciliation may be accomplished in real time or in a batch process. The method, system, and program
10 product further provide notification when the feed records and billing records are out of balance. A further aspect of the method, system, and program product is the capability to scan the various log files on a random or scheduled or as-needed basis, examining the log files for missing, duplicate, and corrupted records, and generate balance reports and out of balance notifications, as necessary.

15

Aspects of the method, system, and program product include a mechanism to uniquely identify feeders, target billing systems, and even individual records. A further aspect of the method, system, and program product is the capability of handling both batch and real time feeds independent of the data format or communications protocol. A still further
20 aspect of the method, system, and program product described herein is the handling of multiple feeders, multiple independent hops, and multiple target billing systems. A still further aspect of the method, system, and program product described herein is reporting capability to show in-balance and out-of-balance conditions.

25 According to the method, system, and program product described herein, an administrator is notified of any error condition, including information for problem determination. The information may include a list of records sent by the various feeders and not received by the billing system, as well as information received by the billing system but neither

END920030154US1

identified to or associated to a feeder. This allows determination as to where the records were lost or duplicated.

FIGURE 2 is a simplified flow chart showing data collection, data reconciliation, and charge data consolidation (that is, billing).

Record reconciliation depends on effective and accurate logging. In this regard, each record sent by the feeder must contain certain tracking information, and each subsystem must log, that is, write a log entry, for each record processed. Each log entry must contain the necessary tracking information.

Record reconciliation in the reconciliation subsystem uses the various logs to reconcile the feeds and for problem determination and reporting. While the most efficient way to maintain these logs is through a database or a system of databases, the method, system, and program product is independent over whether the log is maintained in one or more databases, or in files, or by another mechanism.

The following information should be in each log entry:

1. A unique system identifier, that is, a system identifier that is unique within the scope of the infrastructure being reconciled.
2. A target system identifier, that is, a system identifier that is unique within the scope of the infrastructure being reconciled.
3. A unique record identifier, that is a record identifier that is unique within the scope of the system.
4. The time and date ("time stamp") for the record being logged.
5. The audit scope identifier, that is, the time and date of the next scheduled reconciliation as well as other audit data.

END920030154US1

6. A numeric field indicating the quantity of the resource consumed by the customer.
This field may also contain various product or resource identifiers.
7. A reconciliation time and date to be filled in by the reconciliation system.

5 With respect to the unique system identifier, each subsystem in the system must be uniquely identified. For example, each hardware subsystem in the system could have a unique identifier, with a unique “application instance” identifier for each feeder on the hardware system. These are referred to herein as a SystemID and AppInstanceID, respectively.

10

Identification of the target billing system is essential. This is especially true where the billing system has more than one component or more than one target billing system. This is especially the case with electric power where one transmission system or one local distribution system serves a plurality of generation companies, for example, with
15 different billing structures. This is also the case for the products and byproducts of a modern, integrated petroleum refinery, where a range of products from heavy hydrocarbons (asphalt, tar) through liquid products (as gasoline, diesel fuel, kerosene) to gaseous hydrocarbons (methane, ethane) are either transferred internally (with an associated inter-divisional transfer cost), or sold, for specific industrial or commercial
20 segments.

The target billing systems can be identified in several ways. For example, the target billing system name could be explicitly specified by the feed in the record created by the feeder to be sent to the various target billing systems. Alternatively, the record could
25 contain some verb that indicates to a processing system which billing target system to use.

END920030154US1

According to a still further alternative, there may exist a one to one mapping of specific products to specific target billing systems.

Each record requires a unique identifier. This unique identifier is referred to herein as a Request Reference, "RequestRef". The RequestRef must be unique within the set of all records created by a specific feeder in a specific and configurable period of time. The time period may be arbitrarily or algorithmically set, and is related to how long the owner of the system wishes to retain records for audit, reconciliation, or problem determination purposes. Once a set of records are discarded the RequestRefs of those records may be reused, however, a feeder may not create a new record with a RequestRef that is currently stored in any log.

A further record identifier is the unique key of a record. Each record in the system has a unique key. In one embodiment the unique key may be obtained by concatenating the SystemID, AppInstance, and RefRequest. This unique key would identify a record to a feeder system, a particular feeder, a record from that particular feeder on the particular feeder system. The routines for processing one or both of the log databases and log file processing routines should be programmed to never allow any subsystem in the feeder – reconciliation – billing system to store a new record that has a unique key, including SystemID, AppInstance, and RefRequest that is already contained in any local log. Any attempt to assign the same unique key to another, new, record should generate an error. The error should either prompt the offending feeder to generate a new RequestRef, if the feeder is the source of the error, or return an error message to the feeder that generated the record.

A quantity field is part of a feeder record intended for a billing subsystem. The "quantity" is a numerical value that indicates how much of a resource a customer has consumed. In order to ensure that quantity fields are not altered along the path to the billing subsystem,

END920030154US1

the quantity values may be used as a “hash” value. If quantity is not used, the feeder may enter a random number to be used in the same manner.

5 A further field is a log time and date field. Each log entry must contain the time and date that the entry was made. This entry is used by the reconciliation subsystem to ensure that records are properly reconciled.

10 A further field in each feeder record is the audit scope field. The audit scope data is used to group records for reconciliation. The audit scope field is a character field that contains the year, month, and day, and optionally time, of the next scheduled reconciliation for that feeder. The character field may be YYYYMMDD (year, month, day) or even YYYYMMDDHHMMSS (year, month, day, hour, minute, second).

15 The audit scope field or audit scope ID date indicator changes as often as the feeder is scheduled for reconciliation. For example, if the feeder is scheduled for weekly reconciliation, then the feeder must change the Audit Scope ID weekly to match the next scheduled reconciliation.

20 A further requirement for reconciliation is individual feeder profiles for each subsystem. The profile data can be stored in databases or other files. The profile information must include:

1. Address and access information. This is both the feeder address on the network (including a virtual network), and such access information as system login, database information, and file names.
- 25 2. The feeder’s local time zone.
3. Notification information for an out of balance condition, that is, e-mail, pager, or cell phone information.
4. Information on where to send reconciliation information.

END920030154US1

5. The path information, that is, the system ID for each “hop” in the path that the data will take from the feeder system, through interim processing points, to the target billing system.

5 The automated reconciliation program will generally run for each feeder will normally run on a scheduled time, typically a scheduled time for each individual feeder or set of feeders. However, automated reconciliation may be performed as needed, that is, to ensure that an anomaly is corrected or a problem fixed, or to reconcile feeds outside the normal operating time. Any system scheduler, such as “cron” on UNIX can be used to
10 schedule reconciliations.

FIGURE 3 is a flow chart of the reconciliation process. The reconciliation starts by logging the session start time and retrieving the last session start time, and retrieving the logs from the feeder and the interim control points. If there is an error at this point, the
15 appropriate notification is issued, the session status is logged, and the reconciliation process ended. If there is no error, the records are compared, checked for unreconciled records, and the report prepared and published. Session status is logged, and the reconciliation process is ended.

20 The automated reconciliation program, as shown in FIGURE 3, is started by a system job scheduler to reconcile records within the audit scope for one or more feeders.

Reconciliation is started when the scheduler passes a feeder system ID to the automated reconciliation program.

25

The automated reconciliation reads the profile information for the identified feeders, and uses the current time and the feeder’s time zone to calculate the current scope ID value to use.

END920030154US1

The automated reconciliation subsystem retrieves reconciliation logs along the path used moving data from the selected feeder to the selected billing subsystem. Only records within the audit scope are retrieved for reconciliation.

5

At this point the automated reconciliation sub system verifies that the records sent from the feeder or feeders balance with the records received at the billing subsystem. This includes that the totals of the quantity fields are correct on both the feeder sub systems and the billing sub system, that is, that they match. If the records do not balance (that is, 10 the numbers of records of the various identification fields do not match) or the quantity fields do not match, the automated reconciliation sub system issues an appropriate notification. This notification is issued in accordance with the information in the profile for the feeder.

15 To be noted is that the automated reconciliation sub system has all of the logs for all of the points in the system, that is, feeder sub systems, nodes, interim nodes, and billing sub systems. This enables the automated reconciliation subsystem to determine where a record was lost or altered.

20 The automated reconciliation sub system also checks for old, unreconciled records in the logs. These are records with audit scope dates older then the audit scope date currently being reconciled. If such old, unreconciled records exist, the automated reconciliation sub system issues a notification according to the feeder profile.

25 The automated reconciliation sub system also checks for the earliest log date on the set of records collected for the audit scope ID. The automated reconciliation sub system uses this date to check if there are any records with earlier log dates that are not marked with a

END920030154US1

reconcile date. If any such records are found, the automated reconciliation system issues a notification according to the information in the profile for this feeder.

5 This condition may arise if the feeder sub system sends records with a different audit scope ID than it should have. This may also arise when an audit scope was out of balance due to an extra record having been sent by the feeder sub system, so that the extra record does not have a reconcile date in the log. In this case the record may be detected by the check for unreconciled records in the next scheduled reconciled. An understanding of the set of events giving rise to the unreconciled or extra record is important for doing
10 problem determination for an out of balance report.

If the audit scope is in balance, the automated reconciliation sub system marks the records that were reconciled in appropriate logs, for example, each log in the path. Alternatively, successfully reconciled records can be pruned.

15

After reconciliation, the automated reconciliation sub system sends a balance report in accordance with the profile data, and also logs its own session status and balance information, e.g., for problem determination purposes.

20 Reconciliation can be scheduled or run on an “as needed” basis. Typically, the schedules for reconciliation are set to complete on a cycle that is consistent with the actual billing system cycle, that is, typically, monthly. This helps out of balance conditions from interfering with the billing system and the invoicing cycle. For a real time system, a best effort can be made to quiesce the real-time input into the billing system for month end
25 processing.

The automated reconciliation sub system can be configured or programmed to produce general and specialized balance reports. By way of exemplification and not limitation,

END920030154US1

such reports can include in and out of balance indications, feeder system and AppInstance IDs, audit scope IDs, date and time reconciliation started and ended, number of record processed, and number of errors detected.

5 For each error the automated reconciliation sub system can provide such data as the error message number; the date and time that reconciled was started and ended; the date and time that the error was detected. For a record that was not found in either the interim point logs or the target billing system log, the sub system can provide a copy of the record, the last log in which it was found, and a trace of the record (typically starting at
10 the feeder system long and extending toward the target billing sub system); a copy of old unreconciled records; nonbalancing fields (quantity, total records); and the logged date/time, audit scope ID, Request Ref, Target System ID and Quantity for unreconciled records.

15 FIGURE 4 illustrates a simplified system with a single feeder to a single billing system. The feeder is associated to a feeder log, and to an interim node with an associated interim log. The feeder side of the system is separated from the billing side by a fire wall. The billing side includes an interim node with an interim node log, and the billing system with an associated billing log.

20

FIGURE 5 illustrates a complex system with multiple feeders to a single billing system. The feeders are associated to feeder logs, and to an interim node with an associated interim log. The feeder side of the system is separated from the billing side by a fire wall. The billing side includes an interim node with an interim node log, and the billing system
25 with an associated billing log.

FIGURE 6 illustrates a still more complex system with multiple feeders to multiple billing systems. The feeders are associated to feeder logs, and to an interim node with an

END920030154US1

associated interim log. The feeder side of the system is separated from the billing side by a fire wall. The billing side includes an interim node with an interim node log, and the billing systems with associated billing logs.

5 FIGURE 7 illustrates a data collection – data reconciliation – billing system applied to an electric power distribution system. The data collection – data reconciliation – billing system has multiple feeders, typically “electric meters” to multiple billing systems. The feeders collect data from the various associated “electric meters” and are associated to feeder logs, and to an interim node with an associated interim log. The feeder side of the
10 system is separated from the billing side by a fire wall. The billing side includes an interim node with an interim node log, and the billing systems with associated billing logs. In a purely deregulated environment the different billing systems could be different generating companies, such as a coal fired company, a gas fired company, a hydroelectric company, and a wind turbine company. Alternatively, the different billing systems could
15 represent different markets or regulatory regimes.

The invention may be implemented, for example, by having the reconciliation embodied in software or under the control of software in a processor or server and executing a sequence of machine-readable instructions, which can also be referred to as code. These
20 instructions may reside in various types of signal-bearing media. In this respect, one aspect of the present invention concerns a program product, comprising a signal-bearing medium or signal-bearing media tangibly embodying a program of machine-readable instructions executable by a digital processing apparatus to perform a method for migrating data from one storage medium to another storage medium by reading data from
25 a source volume on a source data storage device, as a bit image of a logical volume; and writing the data to a target volume on a target data storage device as a bit image of a logical volume.

END920030154US1

This signal-bearing medium may comprise, for example, memory in the servers. The memory in the server may be non-volatile storage, a data disc, or even memory on a vendor server for downloading to one of the servers for installation. Alternatively, the instructions may be embodied in a signal-bearing medium such as the optical data storage disc. The optical disc can be any type of signal bearing disc or disk, for example, a CD-ROM, CD-R, CD-RW, WORM, DVD-R, DVD+R, DVD-RW, or DVD+RW. Additionally, whether contained in a cluster with the server or elsewhere, the instructions may be stored on any of a variety of machine-readable data storage mediums or media, which may include, for example, a "hard drive", a RAID array, a RAMAC, a magnetic data storage diskette (such as a floppy disk), magnetic tape, digital optical tape, RAM, ROM, EPROM, EEPROM, flash memory, magneto-optical storage, paper punch cards, or any other suitable signal-bearing media including transmission media such as digital and/or analog communications links, which may be electrical, optical, and/or wireless. As an example, the machine-readable instructions may comprise software object code, compiled from a language such as "C++".

Additionally, the program code may, for example, be compressed, encrypted, or both, and may include executable files, script files and wizards for installation, as in Zip files and cab files. As used herein the term machine-readable instructions or code residing in or on signal-bearing media include all of the above means of delivery.

While the invention has been described with respect to certain preferred embodiments and exemplifications, it is not intended to limit the scope of the invention thereby, but solely by the claims appended hereto.